Proceedings of the 7th Latin American Conference on Industrial Engineering and Operations Management
Panama City, Panama, August 4-6, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/SA07.20260117
Published: August 4, 2026

Governing the Invisible: A Process, Power, and Practice Framework for Responsible AI in Healthcare

Enitza George
Physician Executive, certified Chief AI Officer; Society of Teachers of Family Medicine (STFM) AI Task Force
and Co-Chair, STFM AI Collaborative; Health Systems Program Ambassador
School of Systems Science and Industrial Engineering, Binghamton University
Binghamton, New York, USA
egeorge15@binghamton.edu

Lujain Alhusan
PhD Student, School of Systems Science and Industrial Engineering
Binghamton University, Binghamton, New York, USA
lalhusan@binghamton.edu

Mohammad T. Khasawneh
Director, School of Systems Science and Industrial Engineering
Binghamton University, Binghamton, New York, USA
mkhasawn@binghamton.edu

Abstract

The rapid adoption of artificial intelligence (AI) in healthcare has outpaced the development of organizational management systems capable of ensuring accountability, transparency, and effective oversight. This paper proposes a systems management framework for enterprise AI governance based on three interdependent dimensions: Process, Power, and Practice. Process establishes the operational structures for AI intake, inventory, risk classification, documentation, and lifecycle monitoring. Power defines governance authority, decision rights, and accountability so that AI decisions are guided by policy rather than hierarchy, urgency, or informal influence. Practice represents the organizational behaviors that sustain transparency and responsible leadership when governance is tested. The framework is anchored in a central commitment: no significant AI system should remain invisible, unowned, undocumented, or unreviewed. Prevailing standards and regulatory instruments specify operational processes in considerable detail while addressing decision authority and organizational culture more sparingly, and recent empirical work documents a resulting responsibility vacuum in which AI monitoring is performed inconsistently and owned by no one. The framework is positioned explicitly against six reference sources, illustrated through an applied example of ambient AI documentation adopted outside formal oversight, and operationalized through a set of board-level assurance indicators. The principal contribution is a governance model that shifts the focus from technical controls to organizational system design, treating authority and culture as governance infrastructure rather than as context.

Keywords

Artificial intelligence governance, systems management, organizational accountability, healthcare management and responsible AI.

© IEOM Society International

1. Introduction

Healthcare organizations increasingly depend on artificial intelligence (AI) systems for their day-to-day operations, yet many of these systems remain organizationally invisible. Some are embedded within enterprise software, others are acquired from vendors, and still others are developed internally. Regardless of origin, AI can shape decisions without clear visibility into who owns it, how it is monitored, or who is accountable when problems arise. This lack of visibility creates a governance problem, because organizations cannot govern what they cannot see, and governance approaches developed for conventional operations are not sufficient to ensure the safe, effective, and accountable use of AI in healthcare (de Vere Hunt et al. 2025).

In adopting AI, healthcare organizations are not simply adopting another technology; they are adopting a new form of organizational intelligence whose effects extend beyond operational performance to clinical decision-making, patient safety, regulatory compliance, and public trust (WHO 2021). Unlike most other industries, healthcare organizations must simultaneously balance scientific evidence, ethical obligations, legal requirements, clinical judgment, and organizational accountability. Consequently, AI requires more than an extension of existing technology governance models; it requires a new organizational approach across the enterprise (Freyer et al. 2025).

Recent literature has advanced important principles for responsible AI, including fairness, transparency, explainability, privacy, and risk management, and increasingly recognizes the need for organizational governance mechanisms (Ong et al. 2025). Comparatively less attention has been given to governance as an organizational capability that integrates operational processes, decision-making authority, and organizational culture into a single coherent system. That gap is the problem this research addresses.

This paper proposes the Process, Power, and Practice (3P) Framework as an integrated organizational architecture for responsible AI governance in healthcare. As illustrated in Figure 1, the framework may be understood as a bridge. In engineering, bridges are judged not by their appearance under ideal conditions but by their ability to withstand the greatest loads they are expected to carry. Likewise, the effectiveness of AI governance is demonstrated not during routine operations, but when healthcare organizations encounter clinical urgency, political influence, financial constraints, operational demands, and rapid technological change.

1.1 Objectives

This research pursues four objectives. The first is to characterize, through structured comparison with prevailing standards and regulatory instruments, which domains of organizational AI governance are specified least fully by existing guidance. The second is to develop an integrated organizational architecture that addresses those domains. The third is to position that architecture explicitly against existing frameworks so that its incremental contribution can be assessed rather than assumed. The fourth is to propose governance assurance indicators that render framework performance observable at the executive and board level. Section 9 returns to each objective and states how it has been met.

2. Literature Review

The rapid adoption of artificial intelligence has prompted growing attention to the governance needed for its responsible use in healthcare. Recent literature has established responsible AI as a multidisciplinary field grounded in ethical principles, technical safeguards, and human oversight (WHO 2021, STFM 2026). Emerging frameworks emphasize multidisciplinary oversight, lifecycle management, continuous monitoring, risk assessment, and organizational accountability, recognizing that governance extends beyond technical validation or regulatory compliance. The review below is organized by source category rather than paper by paper, covering consensus standards, regulatory instruments, professional guidance, and institutional implementation studies.

International standards have formalized several of these expectations. ISO/IEC 23894:2023 guides managing risk throughout the AI lifecycle and is intended for use alongside established enterprise risk management practice (ISO/IEC 2023b), while ISO/IEC 42001:2023 specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (ISO/IEC 2023a). In the United States, the NIST Artificial Intelligence Risk Management Framework offers a voluntary, function-based structure organized around Govern, Map, Measure, and Manage (Tabassi 2023). Together these position AI governance as an enterprise responsibility requiring coordinated leadership across clinical, operational, legal, technical, and administrative domains.

© IEOM Society International

The regulatory environment has matured in parallel. Regulation (EU) 2024/1689 establishes risk-tiered obligations and classifies most AI functioning as, or embedded within, a regulated medical device as high risk, with corresponding requirements for risk management, data governance, logging, and human oversight, and it imposes an explicit AI literacy duty on providers and deployers (European Parliament and Council 2024). In the United States, the HTI-1 final rule introduced transparency and source-attribution requirements for predictive decision support interventions in certified health information technology, together with expectations for risk management practices and real-world testing (ONC 2024). These instruments place obligations principally on developers and vendors; they do not specify how a deploying healthcare organization should organize itself to meet its own accountability obligations.

Professional and coalition guidance has begun to address that gap. The Coalition for Health AI has emphasized organizational structures, multidisciplinary oversight, policy development, and continuous monitoring as essential capabilities (CHAI 2026), and its joint guidance with The Joint Commission sets out principles spanning AI policy, governance structures, local validation, monitoring, and workforce education (Joint Commission and CHAI 2025). Institutional case studies document what implementation requires in practice, including the sustained stakeholder alignment, role definition, and committee review needed to move a single predictive model into routine clinical use (Sendak et al. 2020b), the value of standardized model documentation for clinical end users (Sendak et al. 2020a), and the organizational process of establishing institution-wide responsible use guidelines (Saenz et al. 2024).

Empirical work has begun to document the organizational consequences of this emphasis. Owens et al. (2025) interviewed 21 experts involved in healthcare AI implementation and found that monitoring and maintenance responsibilities were poorly defined, inconsistently performed, and institutionally undervalued, producing what they characterize as a responsibility vacuum sustained by incentives favoring rapid innovation over sustained oversight; few institutions in their sample had invested in structured oversight infrastructure. Hussein et al. (2026) systematically reviewed 35 healthcare AI implementation frameworks published between 2019 and 2024, identified seven governance domains, and concluded that available guidance remains fragmented and frequently assumes institutional capacity that many organizations do not possess. In response, they developed a five-level maturity model, the Healthcare AI Governance Readiness Assessment. Hassan et al. (2025) reach a comparable conclusion in proposing a healthcare-specific governance framework. Read together, these studies suggest that the central difficulty is less a shortage of governance principles than an absence of assigned authority and supporting organizational culture.

Taken together, this body of work provides a strong foundation for trustworthy AI in healthcare. Relatively little of it, however, treats governance itself as an organizational capability integrating standardized processes, assigned decision authority, and leadership behaviors into a coherent enterprise system, or examines whether that system holds when organizational conditions are adverse. The safety science literature offers a precedent: high-reliability organizations sustain performance in demanding environments not through procedure alone but through cultural properties such as preoccupation with failure, reluctance to simplify, and deference to expertise rather than hierarchy (Weick and Sutcliffe 2015). The framework proposed here extends that insight to AI governance.

3. Methods

This paper presents a conceptual framework rather than an empirical study. The source corpus comprises peer-reviewed literature on AI governance, oversight, and accountability in healthcare published between 2020 and 2026; consensus standards and voluntary frameworks, including ISO/IEC 42001, ISO/IEC 23894, the NIST AI Risk Management Framework, and WHO guidance; and regulatory instruments and professional society guidance, including Regulation (EU) 2024/1689, the ONC HTI-1 final rule, the combined CHAI and Joint Commission guidance, and the STFM AI Playbook. Documents were assessed in their published form as of January 2026. They were selected for authority and relevance to enterprise-level deployment rather than through an exhaustive systematic search, a limitation revisited in Section 8.

The study follows the theory synthesis approach to conceptual research described by Jaakkola (2020), in which concepts drawn from multiple literatures are integrated into a framework offering an explanation not available from any single source. This design was selected for two reasons. The governance requirements relevant to healthcare AI are distributed across consensus standards, binding regulation, professional guidance, and safety science, so no single body of work can be extended in isolation; and the field lacks the shared organizational vocabulary that systematic empirical comparison would presuppose. Conceptual synthesis is appropriate at this stage of a field's development, preceding rather than substituting for empirical evaluation, and the empirical programme this framework enables is specified in Section 8. Applying that approach, the 3P Framework was developed as follows. Governance requirements described across the source corpus were extracted and grouped by organizational function, distinguishing

© IEOM Society International

requirements that specify an activity to be performed, requirements that specify an authority to be assigned, and requirements that presuppose a behavioral or cultural precondition. These groupings were consolidated through iterative thematic synthesis into three dimensions corresponding to three organizing questions: what must be done, who decides, and how the organization behaves. The resulting framework was stress-tested against recognized governance failure modes, including undisclosed AI use, vendor-driven model change after approval, and pressure to bypass review under clinical or financial urgency. The comparison against individual source documents reported in Table 1 is qualitative; systematic instrument-based scoring of framework coverage was outside the scope of this study and is identified in Section 8 as the principal direction for subsequent work.

4. The Process–Power–Practice (3P) Framework

The 3P Framework conceptualizes responsible AI governance through three interdependent organizational dimensions. Like the structural elements of a bridge, Process, Power, and Practice are distinct yet interdependent, and weakness in any one reduces the integrity of the whole. Figure 1 presents the framework as an integrated architecture, showing the three dimensions supporting a governance deck that spans clinical care, operations, and research, with the competing pressures that test governance represented beneath. The figure is organized this way deliberately: the pillars are drawn as load-bearing rather than decorative, because the central claim of the framework is that governance fails when any one dimension is absent, not merely that all three are desirable.

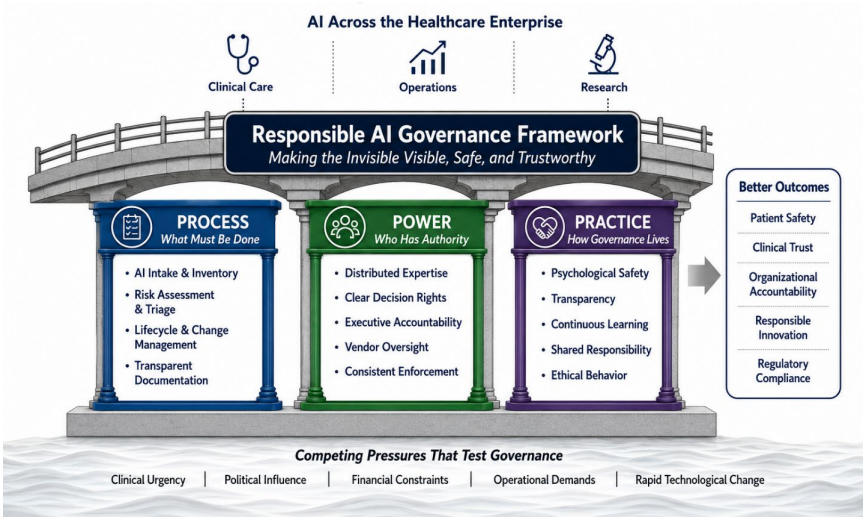


Figure 1. The Process, Power, and Practice (3P) Framework as an integrated governance architecture, in which three complementary dimensions work together to withstand the organizational pressures commonly encountered in healthcare. The illustration was created by the author using generative artificial intelligence and subsequently refined for accuracy and presentation.

4.1 Process

Process defines what must be done to govern AI consistently throughout its lifecycle. Effective governance cannot rely on ad hoc decisions, individual expertise, or isolated technical reviews. It depends instead on repeatable workflows that introduce, monitor, and manage AI consistently according to organizational policy and risk. Four foundational processes are essential.

4.1.1 AI intake, inventory, and risk-based classification

Governance begins by making AI visible, because organizations cannot govern systems that are unknown or poorly understood. This requires a standardized intake process and a comprehensive enterprise inventory covering internally developed models, commercial applications, AI embedded within enterprise software, and clinician-selected tools that operate outside formal oversight, such as ambient documentation applications on personal mobile devices or consumer-grade tools used informally for clinical decision support. Inventories typically integrate procurement records, vendor disclosures, enterprise technology inventories, departmental reporting, and periodic workforce surveys. Because commercial products evolve continuously through software updates, model retraining, and expanded

© IEOM Society International

functionality, the inventory should also capture vendor-driven changes that may alter a system's capabilities or risk profile.

Each identified system should then undergo standardized risk classification against explicit criteria. Although thresholds vary by organizational context, AI should generally be treated as high risk when it influences clinical decisions, patient safety, regulatory compliance, financial reporting, workforce decisions, or other activities carrying significant legal, operational, financial, or reputational consequences. Where an organization operates in a jurisdiction with statutory risk tiers, internal criteria should be reconciled with those tiers rather than maintained in parallel (European Parliament and Council 2024). The purpose of classification is to apply the greatest governance where the greatest risk exists, so that higher-risk applications receive multidisciplinary review and executive oversight while lower-risk applications proceed through streamlined pathways. Because AI evolves, classification must be reassessed whenever functionality, intended use, deployment environment, data sources, or vendor-provided updates change materially.

4.1.2 Documentation: AI system profiles and stakeholder impact assessments

Visibility and classification are sustained through documentation, and two artifacts are central. An AI system profile is a versioned record maintained for every inventoried system, summarizing intended use and target population, model type and provenance, training and validation data, performance in the local deployment environment, known limitations and contraindications, human oversight requirements, the monitoring plan, the assigned owner, and the date of most recent review. The profile is the operational reference for downstream review, monitoring, and audit, extending to institutional governance the standardized model documentation concept developed for clinical end users (Sendak et al. 2020a). Versioning matters: a vendor update or model retraining should generate a new record rather than overwrite the prior one, so that the governance history of a system remains reconstructable.

A stakeholder impact assessment complements the profile by documenting who is affected and how. For higher-risk applications it should identify affected patient populations, the clinicians and staff whose workflow or scope of judgment changes, anticipated benefits, foreseeable harms including differential performance across patient subgroups, implications for workforce roles, and the mitigations and monitoring commitments adopted in response. Completing this assessment before deployment converts equity and safety considerations from aspirational commitments into documented, reviewable governance decisions.

4.1.3 Continuous performance and risk monitoring

AI systems should be monitored throughout operational life to confirm they continue to perform safely, effectively, and as intended. Monitoring should address model drift, declining clinical or operational performance, safety events, algorithmic bias, cybersecurity vulnerabilities, regulatory change, and user feedback. Each inventoried system should carry a defined monitoring plan specifying the metrics tracked, the review frequency, the thresholds that trigger escalation, and the individual accountable for review. Continuous monitoring allows organizations to detect problems early, act promptly, and sustain trust in AI-enabled decisions.

4.1.4 Lifecycle and change management

Governance begins before acquisition and continues through the full lifecycle: ideation, evaluation, procurement, implementation, operational use, monitoring, change management, and retirement. It starts by determining whether AI is the appropriate solution at all, after which candidate solutions should be evaluated for intended use, expected benefits, risks, evidence of validation, regulatory status, interoperability, and organizational fit. Procurement is the highest-leverage stage, because it is where governance expectations become contractually enforceable: requests for proposals and vendor agreements should define expectations for transparency, validation support, notification of significant model or software changes, cybersecurity, regulatory compliance, and mechanisms for suspending functionality when necessary. Although responsibilities may be shared contractually, accountability for safe and effective use remains with the healthcare organization.

Implementation should integrate AI into clinical and operational workflows through local validation, technical integration, user training, workflow redesign, and organizational change management before routine use. Because AI is dynamic, governance cannot end at initial approval. Oversight continues through two distinct mechanisms. Periodic review reassesses performance, safety, fairness, cybersecurity, compliance, and organizational alignment at predefined intervals regardless of whether concerns have arisen. Event-driven review is triggered by material change, including new model versions, expanded intended use, significant software updates, changes in or addition of data sources,

© IEOM Society International

deployment in new clinical environments, and vendor-driven modification. Governance should also include a structured end-of-life process covering replacement, retirement, decommissioning, documentation retention, and organizational learning, so that institutional knowledge is preserved and residual risk addressed rather than inherited silently by whatever system follows.

4.2 Power

Power defines the organizational authority required to govern AI effectively, addressing three questions: who has authority to decide, how decisions are documented and reviewed, and how governance is protected when competing pressures challenge independent judgment. Underlying all three is a single principle: AI may inform decisions, but accountability never transfers from humans to the algorithm. This principle anchors the dimension by keeping authority, responsibility, and accountability explicitly assigned to human leaders and governance bodies regardless of the sophistication or autonomy of the system.

4.2.1 Who has the authority to make AI governance decisions

AI governance requires expertise from across the enterprise, because AI increasingly touches every major organizational function. Governance should include representatives from clinical operations, nursing, quality, patient safety, information technology, data science, legal counsel, compliance, privacy, cybersecurity, finance, human resources, executive leadership, and the board. Governance information should be tailored to decision-making responsibility, with boards requiring strategic oversight and operational teams progressively greater technical detail. Multidisciplinary participation alone does not produce effective governance: organizations must explicitly identify who approves new applications, accepts organizational risk, authorizes deployment, requires corrective action, and determines when a system is modified or retired. Expertise should be distributed across the organization; accountability should not.

Organizations should designate an executive leader, such as a Chief AI Officer or another appropriately empowered executive, to coordinate governance enterprise-wide, provide strategic oversight, ensure consistent process implementation, and manage enterprise AI risk. That executive should also hold pause authority: the ability to temporarily suspend a system when credible concerns arise regarding patient safety, privacy, cybersecurity, regulatory compliance, or equity, allowing investigation and corrective action before operation resumes. The ability to pause an AI system is not evidence that governance has failed; it is evidence that governance is functioning as intended. Pause authority is among the organizational controls least consistently addressed by prevailing standards, and healthcare organizations should therefore expect to establish it explicitly rather than inherit it from a framework.

4.2.2 How decisions are documented, justified, and reviewed

Governance decisions should be formally documented so that authority is exercised transparently and remains accountable over time. Organizations should maintain a centralized repository of governance decisions, including risk assessments, approvals, deployment authorizations, exceptions, mitigation plans, monitoring results, and decisions to modify or retire systems. Recording both the decision and its rationale creates an auditable record supporting organizational learning, leadership continuity, regulatory compliance, and periodic review.

4.2.3 How governance is protected under competing pressures

Responsible governance requires that decisions remain evidence-based rather than pressure-based. Decisions on deployment, continued use, or retirement should rest on objective evidence rather than organizational preference or external pressure, and governance bodies should draw where feasible on independent validation studies, third-party audits, external benchmarks, or expert review to strengthen high-impact decisions. Governance authority may be challenged not only by clinical or operational urgency but by financial interests, external sponsorship, philanthropic donations, research funding, vendor relationships, and institutional investments that create actual or perceived conflicts of interest. Transparent disclosure, recusal where appropriate, and reliance on objective evidence preserve decision integrity under these conditions. Conflict-of-interest handling and the protection of governance independence receive comparatively little attention in prevailing AI governance guidance, which is why they must be established deliberately at the organizational level.

4.3 Practice

Practice defines the organizational behaviors and cultural conditions that allow governance to function in everyday operations. Effective governance extends beyond policies and structures; it requires behaviors that promote

© IEOM Society International

transparency, accountability, continuous learning, and shared responsibility. Without these conditions, governance becomes a compliance exercise rather than an organizational capability.

4.3.1 Psychological safety and transparency

Governance depends on people being willing to make AI visible. This requires a culture in which employees can disclose AI use, report concerns, admit mistakes, and identify unintended consequences without fear of retaliation. Governance cannot manage systems that remain invisible because individuals avoid disclosure, so psychological safety is a governance prerequisite rather than a human resources concern. Organizations should frame AI inventories, risk reporting, and governance reviews as collaborative learning rather than punitive oversight; an incomplete AI inventory is often a cultural failure before it is a governance failure.

Everyone who interacts with AI should feel empowered to raise concerns, including clinicians, staff, administrators, contractors, trainees, patients, and caregivers, and organizations should establish trusted channels for reporting safety issues, bias, privacy concerns, cybersecurity vulnerabilities, and inappropriate use. Reporting should extend beyond adverse events to near misses, close calls, unexpected outputs, and good catches, which frequently provide the earliest indications of emerging vulnerability and create opportunities for proactive intervention. Emphasizing learning from both failures and prevented failures moves the organization from reactive incident response toward the high-reliability posture characterized by preoccupation with failure, sensitivity to operations, and resilience (Weick and Sutcliffe 2015).

4.3.2 Continuous learning and AI literacy

Because AI technologies change continuously, governance cannot rely on static policies or one-time training. Responsible governance requires ongoing organizational learning enabling clinicians, administrators, technical teams, and executives to understand emerging capabilities, associated risks, regulatory expectations, and evolving practice. AI literacy should be delivered through an enterprise-wide education strategy led by the executive responsible for governance, with human resources integrating AI competencies into onboarding, continuing professional development, leadership education, and competency assessment while subject-matter experts supply role-specific content. Department leaders play a critical role by monitoring workforce competency within their areas, identifying needs for upskilling, reinforcing appropriate use, and recognizing emerging risks such as automation bias, overreliance on AI recommendations, and erosion of human oversight. Continuous literacy therefore functions not only as an educational programme but as an early governance mechanism that detects behavioral risk before it produces patient harm or operational failure.

4.3.3 A culture of shared responsibility

Responsible AI governance cannot reside solely within a single department, committee, or executive office. Employees throughout the organization should recognize their role in ensuring safe, ethical, and effective use. Questions such as whether an application has been registered, appropriately reviewed, adequately monitored, or responsibly implemented should become routine considerations embedded in everyday clinical and operational workflow. As these behaviors normalize, transparency, accountability, and responsible use evolve from compliance requirements into shared organizational values. Practice transforms governance from structure into culture, and without trust, transparency, continuous learning, and shared responsibility, even the most sophisticated governance structures are unlikely to achieve their intended purpose.

4.4 Relationship to existing frameworks and standards

The coverage scores establish how much each framework specifies, but not what each contributes. Table 1 provides that qualitative complement, stating for each source its principal contribution, the aspect it specifies less fully, and the 3P dimension that extends it. It clarifies that the 3P Framework complements rather than replaces existing guidance: the standards specify what responsible AI requires, while the framework proposed here addresses how an organization should be structured so those requirements are met under pressure.

© IEOM Society International

Table 1. Qualitative relationship of the 3P Framework to existing AI governance frameworks, standards and guidance.

Framework or instrument	Primary orientation and contribution	Aspect specified less fully	Corresponding 3P emphasis
NIST AI RMF (Tabassi 2023)	Voluntary, function-based structure (Govern, Map, Measure, Manage) for identifying and managing AI risk across the lifecycle.	Sector-specific decision rights; how governance behaves when organizational pressure conflicts with the assessed risk.	Power (decision rights, pause authority) and Practice (behavior under pressure).
ISO/IEC 23894 (ISO/IEC 2023b)	Guidance for integrating AI risk management into established enterprise risk management practice.	Cultural preconditions for disclosure; how unregistered or informally adopted AI enters the risk process at all.	Practice (psychological safety, disclosure) and Process (intake and inventory).
ISO/IEC 42001 (ISO/IEC 2023a)	Certifiable management system requirements for establishing, maintaining and improving an AIMS.	Healthcare-specific clinical and safety context; pause authority; behavioral conditions determining whether the system is used as designed.	Process (healthcare-specific workflows) and Practice (culture sustaining the system).
EU AI Act (Eur. Parl. and Council 2024); ONC HTI-1 (ONC 2024)	Binding risk-tiered obligations, transparency, source-attribution and AI literacy requirements.	Obligations fall principally on developers and vendors rather than on deploying provider organizations.	Process (classification aligned to statutory tiers) and Power (vendor oversight, retained accountability).
CHAI (2026); Joint Commission and CHAI (2025)	Practical, healthcare-specific recommendations on governance structures, local validation, monitoring and education.	Integration of the three domains into a single architecture, and explicit treatment of governance resilience.	The integrated 3P architecture and the assurance indicators in Table 2.
3P Framework (this paper)	Organizational architecture integrating operational process, decision authority and culture into one enterprise capability.	Prospective empirical validation of adoption outcomes (see Section 8).	Process, Power and Practice as interdependent dimensions.

4.5 Governance assurance indicators

A framework becomes a management system only when its performance can be observed. Table 2 specifies an indicator set for use by a governance committee or board, organized so that each indicator can be implemented consistently rather than interpreted ad hoc. Every indicator is defined along four attributes: the computation that produces it, the organizational data source from which that computation draws, the cadence at which it is reported, and the evaluation criterion against which the result is judged.

Three measure types are used, and they are read differently. Coverage proportions express completeness against a denominator of governed systems and apply where the governance expectation is universal, so their target is unambiguous. Elapsed-time measures capture responsiveness, which is a governance property rather than an administrative one, because a process that is slow is a process that will be bypassed. Event counts capture the exercise of authority and the surfacing of concerns; these are interpreted directionally rather than against a fixed target, since neither a high nor a low count is unambiguously favorable. A sustained count of zero pause events, for instance, may indicate that no system has warranted suspension or that the authority is not usable in practice, and distinguishing the two requires inspection rather than a threshold.

Because defensible numerical benchmarks depend on organizational scale, case mix, and regulatory exposure, and because no validated benchmarks yet exist for these constructs, the criteria in Table 2 state the direction of intended performance and the conditions that should trigger escalation rather than prescribing values. Organizations adopting the set should establish their own thresholds during the first reporting cycle and treat those initial values as a baseline for longitudinal comparison rather than as external standards.

© IEOM Society International

Table 2. Governance assurance indicator set, showing measurement approach, data source, reporting cadence and evaluation criterion for each indicator.

Dimension	Indicator	Measurement approach and data source	Cadence	Evaluation criterion
Process	Inventory completeness	Inventoried systems holding a complete, current profile and a named owner, divided by all inventoried systems; drawn from the enterprise AI inventory.	Quarterly	Target of full coverage for high-risk systems; any shortfall reported to the board with a remediation date.
Process	Monitoring coverage	High-risk systems reporting monitoring results within their defined interval, divided by all high-risk systems; drawn from monitoring plan records.	Quarterly	Target of full coverage; a lapsed interval triggers review of continued deployment authorization.
Process	Decision responsiveness	Median elapsed days from intake submission to documented decision, stratified by risk tier; drawn from the intake log.	Quarterly	Organization-set service level per tier; a sustained increase is treated as a bypass risk signal, not an administrative delay.
Power	Decision traceability	Governance decisions recording both a rationale and a named accountable decision maker, divided by all decisions in the period; drawn from the decision repository.	Semi-annual	Target of full coverage; an active exception carrying no expiration date is treated as a control failure.
Power	Exercise of independent authority	Count of pause events with median time to resolution; count of conflict-of-interest disclosures and recusals; drawn from the executive governance log.	Annual	Read directionally rather than against a target; a sustained count of zero warrants inspection of whether the authority is usable in practice.
Practice	Disclosure route	Systems entering the inventory through voluntary disclosure, divided by all newly identified systems, with the remainder attributed to audit, incident or procurement; drawn from intake records coded by discovery route.	Annual	A rising proportion indicates improving transparency; a falling proportion is an early warning of cultural deterioration.
Practice	Competency and safety to speak	Role-specific AI literacy completion rate; proportion of favorable responses on survey items concerning comfort reporting AI concerns without reprisal; drawn from the learning management system and periodic staff survey.	Annual	Completion target set locally; survey items tracked longitudinally against the organization's own baseline rather than an external benchmark.

5. Application of the Framework

The following applied example illustrates how the three dimensions interact. It is an illustrative scenario constructed to demonstrate the framework rather than an empirical case study, and the pattern it describes, informally adopted AI reaching clinical use before governance is aware of it, is precisely the pattern that the governance domains identified in Section 2 as least fully specified leave organizations least equipped to handle. A group of primary care physicians in a multi-site ambulatory practice begins using a commercially available ambient documentation application on personal mobile devices to draft clinical notes. Adoption is organic: clinicians report substantial reductions in after-hours documentation burden and share the tool informally with colleagues. The application was never procured by the

© IEOM Society International

organization, was not reviewed by information technology, privacy, or compliance, and appears in no enterprise inventory. Patient audio is transmitted to a vendor whose data retention and model training practices have not been assessed, and drafted notes enter the medical record after clinician review. The tool therefore influences the legal record of care while remaining organizationally invisible.

Under the Process dimension, the periodic workforce survey and departmental reporting mechanisms described in Section 5.1.1 are the only mechanisms likely to surface the tool, because neither procurement nor enterprise technology inventories would capture it. Once identified, intake and classification apply: the system influences the medical record and processes protected health information, so it is classified as high risk, requiring an AI system profile, a stakeholder impact assessment addressing patients whose consent practices were never defined, and a monitoring plan covering transcription accuracy and error patterns across accents, languages, and clinical specialties. Documentation is what makes the previously invisible system reviewable.

Under the Power dimension, the question is who decides what happens next, and on what evidence. The designated executive may exercise pause authority while privacy and security review proceeds, without that pause being treated as punitive toward the clinicians involved. Decision rights determine who may authorize continued use, who accepts residual risk, and who negotiates the contractual terms, including data retention, model training on patient audio, notification of model changes, and suspension rights, that were never established because the tool was never procured. The example also demonstrates the limits of contractual allocation: the vendor may assume defined responsibilities, but accountability for the integrity of the medical record remains with the healthcare organization.

Under the Practice dimension, the governing question is whether clinicians disclosed the tool voluntarily or whether it was discovered through audit or complaint. In an organization with psychological safety, clinicians raise the tool themselves, often together with the workflow problem that motivated its adoption. In an organization without it, disclosure is deferred, adoption spreads, and governance learns of the system only after an adverse event. The organizational response also shapes future disclosure: responding punitively to good-faith adoption of a tool that addressed a real documentation burden reliably reduces the disclosure the organization will receive next time.

The example illustrates the central claim of the framework. Process alone would have produced an inventory entry but no authority to act. Power alone would have produced authority with nothing visible to exercise it on. Practice alone would have produced disclosure with no structured path for handling it. The governance outcome depends on all three functioning together, which is why the comparative under-specification of two of the three dimensions is consequential rather than merely descriptive.

5.1 Implementation sequence

Adopting the framework in full is not feasible as a single initiative, and the order of adoption matters because the three dimensions are not independent. Three stages are proposed. The first establishes visibility, implementing intake, the enterprise inventory, risk classification, and system profiles, and delivering a risk-tiered register of AI in use. The second assigns authority, designating the accountable executive, defining decision rights, standing up the decision repository, and securing pause authority both in internal policy and in vendor contracts at renewal. The third builds sustaining practice, introducing role-specific AI literacy, non-punitive reporting channels, near-miss capture, and board reporting against the indicators in Table 2.

The sequence is cumulative rather than optional, and two common failure patterns follow from departing from it. Organizations frequently attempt the third stage first, in the form of principles and training, because it is the least disruptive; the framework predicts this will not hold, because disclosure without a route to review produces reports that no one is empowered to act on, and the resulting inaction teaches staff that disclosure is futile. Conversely, organizations that complete the first stage alone produce an inventory that ages, because nothing obliges anyone to maintain it. Visibility without authority decays, and authority without culture is never invoked.

6. Discussion: Governing Under Pressure

The governance domains that prevailing guidance addresses least fully, decision authority and organizational culture, are precisely the domains that determine whether governance holds when it is tested. This is the central concern of the framework: the condition under which governance most often fails is the condition least addressed by the guidance intended to prevent that failure. Engineers do not judge a bridge by how it performs on a calm day; they ask whether it withstands maximum load, severe weather, and unexpected stress without failure. Governance should be evaluated the same way. Its effectiveness is demonstrated not under routine conditions but when organizations face clinical urgency, political influence, financial constraint, operational disruption, or rapid technological change. Organizations

© IEOM Society International

rarely struggle to make ethical decisions when time is abundant and consensus exists; the true test is whether governance enables the right decision when the right decision is difficult.

6.1 Process under pressure

Under pressure, governance succeeds only if it remains practical. Processes that are slow, overly complex, or disconnected from operational reality encourage individuals to bypass oversight in favor of expediency. Effective governance therefore requires clear criteria for disclosure, rapid risk triage, defined escalation pathways, stakeholder impact assessments, and transparent documentation of decisions and exceptions. When governance is predictable, proportional, and responsive, urgency becomes a reason to engage governance rather than to avoid it. The right course of action should remain the easiest course of action even under pressure, which is why median time from intake to decision appears among the assurance indicators in Table 2 as a first-class governance metric rather than an administrative statistic.

6.2 Power under pressure

Under pressure, authority must remain stronger than influence. Responsible governance depends on decision structures in which accountability outweighs hierarchy, politics, or operational convenience, sustained through executive accountability, multidisciplinary review, documented exceptions, and consistent leadership support. Governance must also remain independent of informal organizational influence: personal relationships, executive sponsorship, historical loyalties, and political alliances should neither exempt AI initiatives from oversight nor alter governance outcomes. Informal influence exists in every organization, but responsible governance requires that formal decision rights consistently prevail over informal power structures.

The same principle extends beyond the organization. Healthcare increasingly depends on third-party AI vendors whose products evolve through software updates, model retraining, and expanding functionality outside the organization's direct control. Governance therefore extends to vendor oversight through contractual expectations, transparency requirements, validation of significant system changes, and the authority to suspend or restrict functionality when necessary. Contracts may allocate responsibilities, but they do not transfer accountability. Governance authority may likewise be challenged by financial interests, external sponsorship, philanthropic donations, research funding, and institutional investments that create actual or perceived conflicts of interest. Transparent disclosure, recusal where appropriate, and reliance on objective evidence preserve the integrity of decisions under these conditions. Because prevailing guidance says little about conflict of interest in AI decision-making, organizations cannot expect existing frameworks to supply these safeguards.

6.3 Practice under pressure

Under pressure, culture determines whether governance is followed or abandoned. Even the strongest processes and decision structures cannot succeed without an environment that values transparency and psychological safety. Individuals must feel safe disclosing AI use, reporting uncertainty, raising ethical concerns, and seeking guidance without fear of blame or retaliation. Organizations that discourage transparency inadvertently encourage invisible AI adoption, undermining governance before formal oversight can occur, while organizations that reward openness create conditions for governance to become collaborative rather than punitive. Leaders reinforce this culture by modeling ethical decision-making, encouraging respectful challenge regardless of hierarchy, and demonstrating that patient safety and transparency remain priorities during periods of uncertainty. Psychological safety is therefore not simply an organizational virtue; it is a governance capability, and its thin treatment in prevailing guidance is a substantive gap rather than a matter of emphasis.

Collectively these observations illustrate a broader principle. Healthcare organizations routinely operate under uncertainty, competing priorities, and time pressure, and under such conditions expecting individuals to consistently do the right thing without organizational support is unrealistic. Governance should instead be designed so that the right action remains possible, visible, and professionally supported when the pressure to do otherwise is greatest. Process provides the path, Power protects the decision, and Practice creates the confidence to follow it.

7. Validation

Validation of a conceptual contribution rests on convergence with independent evidence rather than on replication of a measurement. Three lines of comparison support the framework, and each of the four objectives stated in Section 1.1 is addressed by one or more of them.

© IEOM Society International

First, the framework converges with recent empirical findings. Owens et al. (2025), interviewing 21 experts across healthcare AI implementation, report that monitoring and maintenance responsibilities are poorly defined, inconsistently performed, and institutionally undervalued, and that few organizations have built structured oversight infrastructure. Their responsibility vacuum is, in the vocabulary of this paper, a simultaneous failure of Power and Practice: the work is unassigned, and the culture does not surface its absence. Hussein et al. (2026), reviewing 35 healthcare AI frameworks systematically, similarly conclude that guidance is fragmented and assumes institutional capacity many organizations lack. That two independent studies, one qualitative and one a systematic review, identify the same deficit that this framework is constructed to address constitutes triangulation across method types.

Second, the framework converges with published implementation experience. Sendak et al. (2020b) report that integrating a single sepsis prediction model into routine care required sustained stakeholder alignment, trust building, role definition, and committee review, none of which is specified by the standards examined in Table 1; the effort they describe falls almost entirely within the Power and Practice dimensions. Saenz et al. (2024) similarly report that an academic health system had to construct its own governance structure and use guidelines because available frameworks did not supply them. Third, the framework is consistent with established safety science: Weick and Sutcliffe (2015) argue that reliability in hazardous settings derives from cultural properties rather than procedure alone, which predicts that guidance concentrating on process specification will under-determine real-world performance. The pattern reported by all four sources is the pattern their theory anticipates.

These lines of evidence do not establish that organizations adopting the 3P Framework achieve better outcomes, which would require prospective evaluation. They establish the necessary prior claim on which the framework rests: that decision authority and organizational culture are consistently under-supported by the guidance currently available to healthcare organizations, and that this deficit has observable operational consequences. The framework is also complementary to, rather than competing with, the maturity model proposed by Hussein et al. (2026): that instrument assesses how far an organization has progressed, while the present framework specifies the architecture along which progress is made.

8. Limitations and Future Research

Four limitations qualify this work. First, the 3P Framework is a conceptual contribution and has not been prospectively evaluated; no evidence is offered that organizations adopting it achieve better governance outcomes than those that do not, and the applied example in Section 5 is illustrative rather than observational. Second, the comparison with existing frameworks presented in Table 1 is qualitative and interpretive. It reflects the authors reading of each source rather than a scored instrument applied by independent raters, and readers may reasonably assess particular frameworks differently. Third, the source corpus was selected for authority and relevance rather than assembled systematically, so requirements emphasized in guidance outside that corpus may be under-represented. Fourth, the framework reflects a predominantly United States regulatory and organizational context and assumes resources, including a designated AI executive, a multidisciplinary governance committee, an enterprise inventory, and continuous monitoring capacity, that critical access hospitals, rural practices, federally qualified health centers, and small independent practices may be unable to sustain, a constraint also emphasized by Hussein et al. (2026).

Four directions follow. The most immediate is to convert the qualitative comparison in Table 1 into a formal coverage instrument: deriving governance requirements from the source documents independently of the framework, scoring each source against them with multiple blinded raters, and reporting inter-rater reliability. Doing so would test the central premise of this paper empirically rather than interpretively. Beyond that, priorities include multi-site implementation studies evaluating whether adoption improves measurable governance outcomes; psychometric development of a governance maturity instrument based on the three dimensions and the indicators in Table 2, positioned against existing maturity models; and extension of the framework to agentic AI, where the assumption that a human decision maker remains meaningfully in the loop requires explicit re-examination.

9. Conclusion

Artificial intelligence is becoming increasingly embedded throughout healthcare, often operating in ways that are organizationally invisible while influencing clinical, operational, and administrative decisions. As adoption accelerates, responsible AI governance must evolve from a collection of policies into an organizational capability.

All four objectives stated in Section 1.1 have been met. The first was met in Sections 2 and 4.4, which show through structured comparison that prevailing standards and regulatory instruments specify operational process in considerable detail while addressing decision authority and organizational culture more sparingly, a reading corroborated by

© IEOM Society International

independent empirical work. The second was met in Section 4 through the development of the 3P Framework, which integrates operational workflow, decision authority, and organizational culture into a single enterprise capability directed at the domains identified as least supported. The third was met in Table 1, which states for each source what it contributes, what it specifies less fully, and which dimension of the framework extends it. The fourth was met in Table 2, which proposes assurance indicators making governance performance observable at the executive and board level.

The unique contribution of this research is a shift in the unit of analysis for responsible AI, from technical controls and ethical principles to organizational system design, treating decision authority and organizational culture as governance infrastructure rather than as background context. Like a well-designed bridge, governance proves its integrity under load. The measure of responsible AI governance is not whether it functions under ideal conditions, but whether it continues to hold when the pressure to do otherwise is greatest.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process

During the preparation of this work, the authors used Claude to polish the language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

CHAI, AI governance playbooks, Coalition for Health AI, Available: https://www.chai.org/, Accessed on January 10, 2026.

de Vere Hunt, I. J., Jin, K. X. and Linos, E., A framework for considering the use of generative AI for health, *npj Digital Medicine*, vol. 8, art. 297, 2025.

European Parliament and Council, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), *Official Journal of the European Union*, Available: http://data.europa.eu/eli/reg/2024/1689/oj, July 12, 2024.

Freyer, O. et al., Overcoming regulatory barriers to the implementation of AI agents in healthcare, *Nature Medicine*, vol. 31, pp. 3239-3243, 2025.

Hassan, M., Borycki, E. M. and Kushniruk, A. W., Artificial intelligence governance framework for healthcare, *Healthcare Management Forum*, Available: https://doi.org/10.1177/08404704241291226, 2025.

Hussein, R., Zink, A., Ramadan, B. et al., Advancing healthcare AI governance through a comprehensive maturity model based on systematic review, *npj Digital Medicine*, vol. 9, art. 236, 2026.

ISO/IEC, *ISO/IEC 42001:2023 Information Technology — Artificial Intelligence — Management System*, International Organization for Standardization, Geneva, Switzerland, Available: https://www.iso.org/standard/42001, 2023a.

ISO/IEC, *ISO/IEC 23894:2023 Information Technology — Artificial Intelligence — Guidance on Risk Management*, International Organization for Standardization, Geneva, Switzerland, Available: https://www.iso.org/standard/77304.html, 2023b.

Jaakkola, E., Designing conceptual articles: four approaches, *AMS Review*, vol. 10, no. 1-2, pp. 18-26, 2020.

Joint Commission and CHAI, Guidance on the responsible use of AI in healthcare, The Joint Commission and Coalition for Health AI, Available: https://www.jointcommission.org/, September 17, 2025.

ONC, Health data, technology, and interoperability: certification program updates, algorithm transparency, and information sharing (HTI-1 final rule), *Federal Register*, vol. 89, no. 6, pp. 1192-1438, 2024.

Ong, J. C. L., Ning, Y., Collins, G. S. et al., International partnership for governing generative artificial intelligence models in medicine, *Nature Medicine*, vol. 31, no. 9, pp. 2836-2839, 2025.

Owens, K., Griffen, Z. and Damaraju, L., Managing a "responsibility vacuum" in AI monitoring and governance in healthcare: a qualitative study, *BMC Health Services Research*, vol. 25, art. 1217, 2025.

Saenz, A. D., Centi, A., Ting, D. et al., Establishing responsible use of AI guidelines: a comprehensive case study for healthcare institutions, *npj Digital Medicine*, vol. 7, art. 348, 2024.

Sendak, M. P., Gao, M., Brajer, N. and Balu, S., Presenting machine learning model information to clinical end users with model facts labels, *npj Digital Medicine*, vol. 3, art. 41, 2020a.

Sendak, M. P., Ratliff, W., Sarro, D. et al., Real-world integration of a sepsis deep learning technology into routine clinical care: implementation study, *JMIR Medical Informatics*, vol. 8, no. 7, e15182, 2020b.

STFM, *Building AI Excellence in Primary Care: Artificial Intelligence Playbook*, Society of Teachers of Family Medicine, Available: https://stfm.org/media/w4nlpi0f/ai-playbook.pdf, Accessed on January 10, 2026.

Tabassi, E., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, National Institute of Standards and Technology, U.S. Department of Commerce, Gaithersburg, Maryland, 2023.

© IEOM Society International

Weick, K. E. and Sutcliffe, K. M., *Managing the Unexpected: Sustained Performance in a Complex World*, 3rd Edition, Wiley, Hoboken, New Jersey, 2015.
WHO, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance*, World Health Organization, Geneva, Switzerland, 2021.

Biographies

Dr. Enitza George is a physician executive, board-certified family physician, and health systems leader with more than 30 years of experience in clinical practice, academic medicine, and healthcare transformation. She serves on the Society of Teachers of Family Medicine (STFM) Artificial Intelligence Task Force, is Co-Chair of the STFM AI Collaborative, and is Health Systems Program Ambassador for the School of Systems Science and Industrial Engineering at Binghamton University. Dr. George holds a Master of Science in Artificial Intelligence in Healthcare and is pursuing a PhD in Industrial and Systems Engineering, where her research focuses on AI governance, systems management, and organizational accountability in healthcare.

Lujain Alhusan is a PhD student in the School of Systems Science and Industrial Engineering at Binghamton University, State University of New York. She holds a bachelor's degree in genetic engineering and a Master of Science in bioinformatics. Her research applies artificial intelligence to cancer genomics, including the analysis of high-dimensional genomic and multi-omic data for tumor characterization, biomarker discovery, and treatment response prediction. Her interests extend to the governance of AI in healthcare, particularly the validation requirements, documentation practices, and organizational accountability structures needed to deploy genomic and AI-enabled decision support tools safely in clinical settings.

Dr. Mohammad T. Khasawneh is a SUNY Distinguished Professor and Director of the School of Systems Science and Industrial Engineering at Binghamton University, State University of New York. His research focuses on healthcare systems engineering, operations management, advanced analytics, and data-driven decision-making, with particular emphasis on improving healthcare delivery through systems optimization. He directs the Healthcare Systems Engineering program and has led interdisciplinary initiatives integrating industrial engineering with healthcare transformation. He has authored numerous scholarly publications and received sustained research funding in healthcare systems and operations, advancing the application of systems engineering principles to complex organizational challenges including healthcare quality, patient safety, and operational performance improvement.

© IEOM Society International